

Ercole, Carnevale Giulio; Paolo, Clemente Gian

Article

A Bayesian internal model for reserve risk: an extension of the correlated chain ladder

Risks

Provided in Cooperation with:

MDPI – Multidisciplinary Digital Publishing Institute, Basel

Suggested Citation: Ercole, Carnevale Giulio; Paolo, Clemente Gian (2020) : A Bayesian internal model for reserve risk: an extension of the correlated chain ladder, Risks, ISSN 2227-9091, MDPI, Basel, Vol. 8, Iss. 4, pp. 1-20, <https://doi.org/10.3390/risks8040125>

This Version is available at:

<https://hdl.handle.net/10419/258078>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.


If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>

Article

A Bayesian Internal Model for Reserve Risk: An Extension of the Correlated Chain Ladder

Carnevale Giulio Ercole ¹ and Clemente Gian Paolo ^{2,*} 

¹ PartnerRe, Hardstrasse 301, 8005 Zürich, Switzerland; giulio.ercole.carnevale@gmail.com

² Department of Mathematics for Economic, Financial and Actuarial Sciences, Università Cattolica del Sacro Cuore, Largo Agostino Gemelli 1, 20123 Milan, Italy

* Correspondence: gianpaolo.clemente@unicatt.it

Received: 28 September 2020; Accepted: 9 November 2020; Published: 19 November 2020



Abstract: The goal of this paper was to exploit the Bayesian approach and MCMC procedures to structure an internal model to quantify the reserve risk of a non-life insurer under Solvency II regulation. To this aim, we provide an extension of the Correlated Chain Ladder (CCL) model to the one-year time horizon. In this way, we obtain the predictive distribution of the next year obligations and we are able to assess a capital requirement compliant with Solvency II framework. Numerical results compare the one-year CCL with other traditional approaches, such as Re-Reserving and the Merz and Wüthrich formula. One-year CCL proves to be a legitimate alternative, providing values comparable with the more traditional approaches and more robust and accurate risk estimations, that embed external knowledge not present in the data and allow for a more precise and tailored representation of the risk profile of the insurer.

Keywords: Bayesian models; stochastic claims reserving; claims development results

1. Introduction

The most recent regulatory requirements for European insurers, known as Solvency II, introduced the possibility of tailored risk modeling for many sources of risks borne by each undertaking. As an alternative to the Standard Formula, according to its own technical capabilities and upon approval of the supervisory authority, each insurer can develop its own internal model. For a specific risk, or a particular set of risks, such a partial internal model is supposed to capture the insurer's risk profile in a more consistent way than the Standard Formula, equal for all the market participants. To accomplish this task, actuarial literature produced several models. In this paper, we provide a partial internal model for assessing reserve risk's capital requirement using Bayesian procedures. In the context of tailored risk modeling, a Bayesian approach has indeed a great potential. In fact, Bayesian techniques are able to include in statistical models external knowledge and subjective judgments that are not necessarily deduced by the data, other than the fact that they require the explicit specification of all the assumptions made. If properly used then, these features enable a better description of the risks faced by an insurer. Anyway, the use of techniques with a Bayesian flavor is not something new for actuarial science (see [Klugman \(1992\)](#)) and with respect to claims reserving can be traced back to [Bornhuetter and Ferguson \(1972\)](#). Despite this, Bayesian models often proved to be too complex to be implemented on real world applications as this kind of modeling often leads to mathematically intractable forms. However, this mathematical complexity can be overcome using Markov Chain Monte Carlo (MCMC) simulation techniques that provide empirical approximations when closed formulas are not available. With respect to claims reserving, the first appearance of MCMC procedures is credited to Ntzoufras and Della Portas in 2002 [Ntzoufras and Dellaportas \(2002\)](#),

which has been followed by many other papers (see, e.g., Peters et al. (2017) and Wüthrich (2007) for recent applications in the field of claims reserves).

Concerning claim reserving, one of the most popular deterministic methods is the chain ladder (see, e.g., Friedland (2010) and Hindley (2017) for a review of main deterministic methodologies). As well-known, the ultimate cost of each accident year is predicted using run-off triangles of paid and/or incurred losses and assuming that prior patterns of losses persist also in the future (e.g., in terms of settlement speed or behavior of the case reserve). Several stochastic methods have been proposed in the literature in order to measure a variability of the Chain Ladder methodology (see, e.g., Mack (1993) and Wüthrich and Merz (2007)). In this context, it is worth to be mentioned the Correlated Chain Ladder (see Meyers (2015); Frees et al. (2016)) that exploits the advantages of Bayesian models and allows to model a correlation between accident years. In this framework, we move from the Correlated Chain Ladder (CCL) and we provide an extension of this approach in order to model the claim development result distribution in a framework compliant with Solvency II. On a one-year time-frame, this proposal represents an alternative of two classical approaches¹, Re-Reserving Diers (2009) and Merz and Wüthrich formula Wüthrich et al. (2008), widely used in practice when chain ladder is the underlying deterministic method. However, this approach is not a mere third option. Indeed, Bayesian techniques represent a more refined and sophisticated approach to obtain claims reserves variability, and this is especially true in comparison to bootstrap based algorithms² (see, e.g., Hastie et al. (2009)). In the context of modeling regulatory capital, using a fully specified probabilistic framework, able to integrate external information, provides a more advanced approach than methods that simply fit semi-parametrically the variability on a limited set of observations. Typically, when speed and simplicity are critical, bootstrap estimates could be more appropriate, while when accuracy is important, Bayesian modeling could represent a more consistent choice. With respect to stochastic claim reserving, we think Bayesian techniques can have a role in improving current standards where the accuracy of estimates is critical, like internal modeling practice.

This paper is organized as follows. In Section 2, the original CCL model, introduced in Meyers (2015), is described. In Section 3, we provide the extension of the model in a one-year view. In particular, we adapt the idea of a Bayesian update originally found in Meyers (2017) in order to assess the distribution of claims development result. In Section 4, we develop a case study based on motor third-party liability data. In particular, in Section 4.1, we focus on the assumptions and on the calibration procedure. In Sections 4.2 and 4.3, we provide main results comparing our proposal to classical approaches provided in the literature. We show how one-Year CCL can be a viable alternative to assess a capital requirement compliant with Solvency II regulation. Conclusions follow.

2. The Correlated Chain Ladder

In Bayesian claims reserving, inference is performed by computing a posterior distribution of the parameters. This distribution is derived via Bayes theorem by combining a prior, representative of external knowledge, and observations. To this extent, we need a data structure that eases the procedure

¹ In this framework, the dynamics of the one-year change in P&C insurance reserves estimation has been also studied in Dacorogna et al. (2018) by analyzing the process that leads to the ultimate risk in the case of “fixed-sum” insurance contracts.

² Statistical literature drew many parallels between Bayesian MCMC and Bootstrap procedures. In statistical inference, these algorithms represent two different ways offered respectively by the Bayesian and frequentist paradigms to assess parameter uncertainty (see, e.g., Efron (2011); Hastie et al. (2009)). The Bootstrap algorithm, and especially its non-parametric version, generates many samples from the data, from which is possible to obtain the distribution of a statistic or a parameter, mimicking the Bayesian effect of a posterior distribution. This distribution is produced in a quick and simple way, requiring no probabilistic assumptions, no prior elicitation nor MCMC procedures and can be effectively considered an “approximation of a non-parametric, non-informative posterior distribution” (Hastie et al. 2009).

of prior parameters specification. Concerning reserving models, we refer to a *cross-classified* structure. Given a run-off triangle:

$$D^{(t)} = \{X_{ij} : i = 0, \dots, t; j = 0, \dots, t; i + j \leq t\}$$

where i, j , and t represent the accident year, the development year, and the evaluation period respectively, and being X_{ij} whatever claims figure (typically either incremental or cumulative payments). In general, a cross-classified structure has the following form:

$$X_{ij} = f(\alpha_i, \beta_j)$$

Claims figures are thought as a function of parameters, where α_i acts as a parameter linked to the accident year and β_j as a development parameter.³ This is a potentially superior informative structure that allows us to link the data to phenomena separately related to either the accident or the development year. With respect to Bayesian inference, this structure is clearly best suited to receive prior information about different factors that affect the reserving process.⁴

2.1. Model Specification

The Correlated Chain Ladder, introduced by Meyers (2015), is basically a re-parameterization of the traditional Chain Ladder model, that exploits exactly a cross-classified structure in order to enhance the possibilities of the traditional algorithm and re-interpret it in a Bayesian context. Given a triangle of cumulative payments:

$$D^{(t)} = \{C_{ij} : i = 0, \dots, t; j = 0, \dots, t; i + j \leq t\}$$

data are assumed generated by a log-normal distribution:

$$C_{ij} \sim \text{logNormal}(\mu_{ij}, \sigma_j)$$

for $i = 0, \dots, t$ and $j = 0, \dots, t$, where:

- $\mu_{0j} = \alpha_0 + \beta_j$
- $\mu_{ij} = \alpha_i + \beta_j + \rho \cdot [\log(C_{i-1,j}) - \mu_{i-1,j}]$ for $i > 0$

In this framework, a *level* parameter α_i embeds information about accident year i , β_j is a sort of *new development factor*, that captures information about the development of payments and ρ is a feature that models the correlation between different accident years. This last parameter is one of the most remarkable innovations brought by this model, in fact it relaxes the hypothesis of independence between accident years, one of the main assumptions of several stochastic models related to Chain Ladder (see Mack (1993); Wüthrich et al. (2008)).

2.2. Parameter Specification

At this point we can elicit prior distribution for each parameter of the model. This step is critical (see Berger et al. (2009)), as we have to choose reasonable distributions and values that are consistent with our data. As a general case, we have the following:

$$\alpha_i \sim \text{Normal}(\log(B_i) + [\log(etr_i)], \epsilon_i) \quad \text{for } i = 0, \dots, t \quad (1)$$

³ In general, parameters could be more than two and left up to the needs of the modeler.

⁴ Anyway this structure is not exclusive of Bayesian models as it is applied also in frequentist models, as for instance the Over-Dispersed Poisson bootstrap Model England and Verrall (1999).

where B_i represents known gross premiums related to the accident year i and $[\log(elt_i)]$ is a further random variable representative of the logarithm of the expected loss ratio for each accident year i , while ϵ_i is the precision, rather than the variance, in accordance to the parameterization adopted by most inference engines. It is necessary to further elicit a prior distribution for $\log(elt_i)$ and as a general case we specify a non-informative uniform prior with parameters γ_i and δ_i (as defined in Meyers (2015)):

$$\log(elt_i) \sim \text{Uniform}(\gamma_i, \delta_i). \tag{2}$$

Obviously, the model can be easily adapted by defining different parameterizations or selecting different distributions⁵. By virtue of the parameterization in (1) and (2), the level parameter α_i describes the logarithm of the ultimate cost of the accident year i . This hyperparameter is critical since its mean is defined on the log-space and even a slight variation of the parameters highly affects the final result. As a result of this, we should state a reasonable domain of variation of the parameters that will result in values of the claims reserve consistent with our data.

The *development* parameter β_j is defined as a uniform random variable on a negative support:

$$\beta_j \sim \text{Uniform}(-\eta, 0) \quad \text{for } j = 0, \dots, t - 1 \tag{3}$$

with $\eta \in \mathbb{R}^+$. According to this assumption, the parameter β_j can be interpreted as the portion of the ultimate cost paid up to the development year j and hence, it is strictly depending on the settlement speed of the portfolio. For $j = t$, we set β_j equal to 0, assuming that all claims are settled until the development year t . The approach can be easily extended in case the estimation of a tail is needed. Regarding the *log-variance*, we have:

$$\sigma_j^2 = \sum_{h=j}^t \tau_h \tag{4}$$

where $\tau_h \sim \text{Uniform}(0, 1)$.

As for the correlation parameter, the choice is still left to the modeler. As a general case, it is possible to set a non-informative uniform prior:

$$\rho \sim \text{Uniform}(-1, 1) \tag{5}$$

2.3. Posterior and Predictive Distributions

After model specification and prior elicitation then we use an inference engine, like STAN or JAGS, to generate samples from the posterior distribution. This step is critical as well, as we need a strong sample that we can deem truly representative of the posterior distribution of parameters.

To this extent, we suggest to sample from n different chains that start from n different points, to allow for a sufficiently long *warm-up* phase and to set an appropriate *thin factor*.

At this point we have K parameters sets $\Theta^{(k)}$, each one representative of a reserving scenario:

$$\Theta^{(k)} = \left(\{\alpha_i\}_{i=0}^t, \{\beta_j\}_{j=0}^{t-1}, \{\sigma_j\}_{j=0}^t, \rho \right)_k \quad \text{with } k = 1, \dots, K.$$

From the posterior distribution, we can obtain the predictive distribution for each cell of the lower triangle, and in particular for the last column. This will allow us to obtain the *predictive distribution* of the ultimate cost and thus of the claims reserve. For each k , we sample this ultimate cost from a log-normal:

$$C_{it}^{(k)} \sim \text{logNormal}(\mu_{it}^{(k)}, \sigma_t^{(k)}) \quad \text{for } i = 1, \dots, t. \tag{6}$$

⁵ For instance, different distributions may be chosen for different accident years i .

As usual, we get the k -th determination of the predictive distribution of the claims reserve by the following formula:

$$R^{(k)} = \sum_{i=1}^t C_{it}^{(k)} - \sum_{i=1}^t C_{i,t-i}. \quad (7)$$

Iterating the procedure K times, we have a Bayesian predictive distribution of the claims reserve in a total run-off framework.

3. An Extension in a One-Year Framework

Now we propose a way to adapt the correlated Chain Ladder on the one-year time horizon in order to describe the reserve risk according to Solvency II guidelines. In this context, a solvency capital requirement (SCR) for reserve risk, at the end of time t , can be derived as:

$$SCR^{0.995} = VaR_{0.995} \left(\sum_{i=1}^t P_{i,t-i+1} + R^{D^{(t+1)}} \right) v_1 - R^{D^{(t)}} \quad (8)$$

where $P_{i,t-i+1}$ denotes the incremental payments of calendar year $t + 1$ for claims incurred in the accident year i , $R^{D^{(t)}}$ is the best estimate at time t and $R^{D^{(t+1)}}$ is the best estimate at time $t + 1$ considering only existing claims at time t . Hence, the capital requirement at a confidence level equal to 99.5% is obtained as the difference between the Value at Risk of the distribution of the next year insurer obligations, opportunely discounted with the risk-free discount factor v_1 , and the best estimate⁶ at time t . Being $R^{D^{(t)}}$ a known amount at the valuation date, we can directly focus on the sum of next year payments for claims already incurred and the residual reserve that will be posted after experiencing one more year of run-off.

Adapting a reserving model on a given time horizon means making an assumption about the payments patterns on such a time horizon and then computing a residual reserve in light of this new information. To this extent, the CCL tries to ground this informative update on a Bayesian approach, making the Bayes theorem a starting point for the determination of the residual reserve after one year of run-off. Main idea is to simulate a set of losses given the parameters sets generated with the MCMC procedure and then re-evaluate the probability of the same sets using the Bayes theorem. Each set could be thought as a different scenario for the reserving process and this effectively means to re-weight probabilities of future scenarios in light of latest realizations. The idea of Bayesian update is originally found in Meyers (2017) and Frees et al. (2016). Here, we adapt it in order to provide an algorithm capable of producing a predictive distribution of *the next year obligations* over the one-year time horizon. Such a distribution is compliant with Solvency II guidelines and is suitable to represent an internal model for reserve risk. Our goal then will be modeling the following random variable:

$$Y^{1yr} = \sum_{i=1}^t P_{i,t-i+1} + R^{D^{(t+1)}}. \quad (9)$$

To do this, we perform the following steps:

⁶ According to Article 76 of the Solvency II Directive [European Commission \(2009\)](#), the claims reserve must be equal to the current amount that insurance and reinsurance undertakings would have to pay if they were to transfer their insurance and reinsurance obligations immediately to another insurance or reinsurance undertaking. This definition leads to a claims reserve evaluated as the sum of the best estimate and risk margin. As prescribed by Solvency II, risk margin is not considered in (8) to avoid problems of circularity.

1. The starting point is represented by the K parameters sets $\Theta^{(k)}$ simulated with the MCMC procedure.

$$\Theta^{(k)} = \left(\{\alpha_i\}_{i=0}^t, \{\beta_j\}_{j=0}^{t-1}, \{\sigma_j^2\}_{j=0}^t, \rho \right)_k \quad \text{with } k = 1, \dots, K$$

We rearrange them according to the original CCL parametrization in order to proceed with the simulations. For each k , we have:

$$\begin{aligned} \mu_{0j} &= \alpha_0 + \beta_j && \text{for } i = 0 \text{ and } j = 0, \dots, t \\ \mu_{ij} &= \alpha_i + \beta_j + \rho [\log(C_{i-1,j}) + \mu_{i-1,j}] && \text{for } i = 1, \dots, t \text{ and } j = 0, \dots, t \end{aligned}$$

where, being the cumulative payment $C_{i-1,j}$ not available in lower triangle (i.e., for $i + j > t$), we simulate it from a $\logNormal(\mu_{i-1,j}, \sigma_j)$.

Hence, for each parameters set we have a $t \times t$ matrix $M^k = [\mu_{ij}^k]$ containing the *posterior log-mean* parameters and a vector $\sigma^k = [\sigma_j^k]$ containing *posterior log-variance* parameters.

2. We simulate the sets of losses starting from the elements of M and σ . Following a one-year approach, we only generate, for each k , S different realizations of the next year future cumulative payments $C_{i,t-i+1}^{(s,k)}$ (with $i = 1, \dots, t, k = 1, \dots, K, s = 1, \dots, S$), hence performing a total of $K \cdot S$ simulations. We end up with an array of dimension $K \cdot S$ where each element is a trapezoid $T^{(s,k)}$ composed of the original triangle $D^{(t)}$ and a further simulated diagonal related to payments in the calendar year $t + 1$. In Table 1 we provide a visualization to clarify this object: each column stands for a different parameter set k and each row is representative of the s -th batch of simulations over the parameters sets, i.e., a Bayesian predictive distribution.
3. Next year incremental payments reported in Formula (9) can be easily obtained by transforming the array of simulated cumulative values in an array of incremental amounts with the following formula:

$$P_{i,t-i+1}^{(k,s)} = C_{i,t-i+1}^{(k,s)} - C_{i,t-i} \tag{10}$$

for $i = 0, \dots, t$. In this step, we obtained $K \cdot S$ diagonals of simulated values of the next year payments.

4. Now, we perform the learning update. We use the information generated by the s -th batch of simulations to update the probability of each parameter set. To this extent, we compute the likelihood of the observations for each element of the array. Being $\Phi_L(x|\mu, \sigma)$ the density function of a log-normal random variable, for every element we have:

$$L[T^{(s,k)}|\Theta^{(k)}] = \prod_{C_{ij}^{(s,k)} \in T^{(s,k)}} \Phi_L\left(C_{ij}^{(s,k)}|\mu_{ij}^{(s,k)}, \sigma_j^{(s,k)}\right) \tag{11}$$

Given the sample of K sets, resampling from the sample guarantees that the probability of sampling a specific set is the same for all the sets. In other words, all the sets initially sampled from the MCMC procedure are equally likely. Then, by means of the Bayes theorem, it is possible to obtain the posterior probability of the k -th parameter set given the simulated losses as:

$$P[\Theta^{(k)}|T^{(s,k)}] = \frac{L[T^{(s,k)}|\Theta^{(k)}] \cdot P(\Theta^{(k)})}{\sum_{h=1}^K L[T^{(s,h)}|\Theta^{(h)}] \cdot P(\Theta^{(h)})} = \frac{L[T^{(s,k)}|\Theta^{(k)}]}{\sum_{h=1}^K L[T^{(s,h)}|\Theta^{(h)}]} \tag{12}$$

In other words, for each simulation, we re-evaluate the probability of each set obtaining a posterior probability distribution. If each set is representative of a possible scenario of the reserving process,

we have effectively the posterior distribution of all possible scenarios. By iterating the process we obtain S different posterior distributions for the parameters.

- By virtue of the log-normal assumption, for each parameters' set, the expected cumulative payments at times greater than t can be computed as:

$$E[C_{ij}^{(k)}] = \exp\left\{\mu_{ij}^{(k)} + \frac{\sigma_j^{2(k)}}{2}\right\} \tag{13}$$

for $i = 0, \dots, t, j = 0, \dots, t$ and $i + j > t$. The Best Estimate $R_k^{D(t+1)}$ easily follows.

For each batch of simulations s , we use the posterior distribution computed at step 4 in order to obtain a *post run-off* re-weighted next year Best Estimate. For every s :

$$\hat{R}_s^{D(t+1)} = \sum_{k=1}^K R_k^{D(t+1)} \cdot P[\Theta^{(k)} | T^{(s,k)}] \tag{14}$$

Thus, we are able to obtain a predictive distribution of the expected values of the residual reserve according to Solvency II standards.

- The predictive distribution of claims development results can be derived by assessing the distribution of the obligations at the end of year $t + 1$. For each simulation, we add to each value of the residual reserve distribution the realized diagonal. For each s , we have at disposal K batches of diagonals over all the parameter sets. Thus, we sample a diagonal and we add it to the value calculated at step 6. In this way, we obtain the s -th realization of the next year obligations:

$$Y_s^{1yr} = \left(\sum_{i=1}^t \hat{P}_{i,t-i+1}^{(s)}\right) + \hat{R}_s^{D(t+1)} \tag{15}$$

Again, by iterating this process over all the S batches of simulations, we finally obtain a predictive distribution of future obligations.

Table 1. We provide an immediate visualization of what was described in the step 2. T stands for the trapezoid obtained adding to the original triangle the new diagonal containing simulated payments of first calendar year after the evaluation date. The upper parenthesis locates the simulated diagonals in terms of parameter set k and number of simulations s .

Simulation Array					
	$\Theta^{(1)}$...	$\Theta^{(k)}$...	$\Theta^{(K)}$
1	$T^{(1,1)}$...	$T^{(1,k)}$...	$T^{(1,K)}$
...		
s	$T^{(s,1)}$...	$T^{(s,k)}$...	$T^{(s,K)}$
...		
S	$T^{(S,1)}$...	$T^{(S,k)}$...	$T^{(S,K)}$

We can then use this distribution to compute the reserve risk capital charge according to the Solvency II directive by applying (8), i.e., by subtracting the Best Estimate to the discounted 99.5% quantile.

Previous steps have been described in order to explain the framework of the algorithm we are providing. When coding it, many short-cuts could be taken. For instance, in Table 1 the array is composed of trapezoids. It is not strictly necessary to save the original triangle for each element, not only because of computational speed issues but also because this information will be simplified

when computing the learning update. In general, the algorithm could be implemented not necessarily following the outline we provided.

We provide also a generalization on an n -year time frame, allowing to obtain a predictive distribution for the *next- n -years obligations*. Mathematical details are reported in Appendix A.

4. An Application of the Model

4.1. Dataset and Model Calibration

In this section, we describe the results obtained applying the correlated chain ladder to a 11×11 triangle representative of a *Motor third party liability* portfolio⁷. The triangle has been obtained simulating cumulative payments from a log-normal distribution starting from an observed triangle to preserve the confidentiality of the data. This allows to check if the parameters of the posterior distributions are consistent with the data used.

As described later, the posterior distribution has been obtained using a Hamiltonian Monte Carlo (HMC) procedure. To implement the methodology, we maintained the log-normal hypothesis for the observed cumulative payments in the triangle and we elicited the following prior structure, structured according to our knowledge about the business.

In particular, the log-ultimate cost is approximated by the parameter α_i , defined as follows:

$$\alpha_i \sim \log(B_i) + [\log(elt_i)] + u_i$$

where:

- $\log(B_i)$ is the natural logarithm of the premium earned for the i -th generation.
- $[\log(elt_i)]$ is the natural logarithm of the *expected loss ratio*, for the i -th generation.
- u_i is a random noise defined as follows:

$$\begin{cases} u_i = 0 & \text{for } i = 0, \dots, 3 \\ u_i \sim \text{Uniform}(-0.6, 0.6) & \text{for } i = 4, \dots, 10 \end{cases}$$

We deemed the variability of the loss ratio sufficient, given that oldest accident years are almost closed.

With respect to the *elt* variable, we choose a log-normal distribution centered on a prior loss ratio differentiated by accident year, mainly based on the analyses made by the company for underwriting purposes. Values are specified in Table 2. Since a limited claims pattern is observed for recent accident years, a higher volatility is assumed.

⁷ Cumulative payments are reported in Table A1.

Table 2. We report the priors we elicited for each accident year. The older the accident year, the more evident is the final loss ratio and then the smaller the variability parameter. With respect to years 1 and 2, the variability is negligible as the final development years will have a non-significant effect on the final loss ratio. On the other hand, for the three youngest accident years, we opted for a quite wide distribution.

Years	Prior Loss Ratio	Log-Mean	Log-Std
1	0.73756	−0.30645	0.000005
2	0.7178	−0.3336	0.000005
3	0.731	−0.31531	0.001
4	0.806	−0.2177	0.008
5	0.85	−0.16455	0.025
6	0.685	−0.38037	0.035
7	0.787	−0.24156	0.05
8	0.705	−0.35159	0.08
9	0.72	−0.33054	0.08
10	0.74	−0.30314	0.085

For development parameter β_j , we have:

$$\begin{cases} \beta_j \sim \text{Uniform}(-3, 0) & \text{for } j = 0, \dots, 9 \\ \beta_j = 0 & \text{for } j = 10 \end{cases}$$

With respect to the variability parameter, we recall that σ_j^2 is defined as:

$$\sigma_j^2 = \sum_{h=j}^t \tau_h$$

where we set $\tau_h \sim \text{Beta}(1, 7)$ ⁸.

Finally, we kept the original assumption about the correlation parameter ρ :

$$\rho \sim \text{Uniform}(-1, 1)$$

using a non-informative prior. Having formalized these assumptions, we generated a posterior distribution of 10,000 parameters sets with the help of STAN inference engine. In this sampling, we stuck to the default sampling procedure adopted by STAN: the No-U-Turn Sampler (NUTS), an extension of HMC procedure (for details, see Algorithm 2 in Hoffman and Gelman (2014)). As shown in Hoffman and Gelman (2014), empirically, NUTS performs at least as efficiently as a well tuned standard HMC method and, as shown later in the document, allowed us to obtain a strong sample.

In particular, we simulated 4 chains, each composed of 27,500 iteration, utilizing the first 2500 as warm-up and thinning the chain by a factor of 10.⁹ With this parameterization we traded computational time for accuracy: for each parameter, the *R-hat* statistic is equal to 1, the traceplot shows a well centered random noise, and the *effective sample size* is well above 9000, implying that we have a strong sample, truly representative of the posterior. For each parameter of the posterior, we provide in Table 3 the *effective sample size*¹⁰, a measure of equivalence between dependent and

⁸ As a general rule we advise to check the order of magnitude of parameters implied by the observations in the triangle, in order to feed the model with consistent inputs and speed up convergence.

⁹ With *adapt-delta* = 0.98 and *maximum treedepth* = 13.

¹⁰ Computed by the package *rstan*.

independent samples.

Table 3. The *thin factor* of 10 and the tight parameterization for the HMC procedure ensured the strong sample.

Effective Sample Size							
α_0	-	β_0	9740	σ_0^2	9687	elr_0	10,152
α_1	-	β_1	10,214	σ_1^2	9719	elr_1	9586
α_2	-	β_2	10,130	σ_2^2	9564	elr_2	10,361
α_3	-	β_3	10,010	σ_3^2	9137	elr_3	9244
α_4	10,024	β_4	10,290	σ_4^2	9413	elr_4	10,089
α_5	9902	β_5	9940	σ_5^2	9343	elr_5	10,048
α_6	10,362	β_6	9595	σ_6^2	10,044	elr_6	10,387
α_7	9736	β_7	9875	σ_7^2	9626	elr_7	9767
α_8	9874	β_8	9743	σ_8^2	9773	elr_8	9595
α_9	9936	β_9	9605	σ_9^2	9863	elr_9	9935
α_{10}	10,166	ρ	9747	σ_{10}^2	9872	elr_{10}	10,064

Additionally, we have verified that the parameters used to simulate cumulative claims are in line with the parameters of the posterior distributions selected by the procedure. To give an indication of the behaviour, we report in Appendix B, Tables A2 and A3, the comparison between these values. As expected, there is not much difference between the parameter μ used to simulate and the mean of the a posteriori distribution. We have indeed that the values are very close with an absolute difference higher than 0.5% only for some payments of the accident year 0. This is further confirmation of the successful sampling with the MCMC algorithm. Finally, we show in Table A4 how our prior structure in terms of loss ratio is reflected very well in the posterior loss ratios.

The important feature of the CCL, and more in general of Bayesian models, is that it embeds in itself a *learning procedure*. In fact, the posterior distribution of parameters is a mix between the prior opinions of the modeler and the experience represented by the data. By looking at them we can effectively see what the model *learned* from both the data and our prior. In Figure 1 we show the caterpillar plots for the level, development, log-variance, and *expected loss ratio* parameters, which provide an immediate overview of the learning update. In the top right, we see the log-variance parameters σ^2 . Their behavior is exactly the one we imposed when eliciting the prior structure as the variability is decreasing with respect to the increase in the development year. In the top left we see the level parameters, where we immediately notice that they are almost all centered on zero. On bottom left we have the development parameters, where we see that a posteriori they acquired a meaning consistent with the fact that the model is built on cumulative payments. These parameters represent the negative shift from the log-ultimate cost: for parameters corresponding to the first development year this shift is consistent, for the last is really tiny. Lastly, the plot at the bottom right represents the highest posterior density regions for the *expected loss ratio*. Again, we can immediately see a behavior consistent with our assumptions, where the domain of variation is wider for the youngest generation and really tiny, if not *non-existing*, for the older ones.




Figure 1. We show the caterpillar plots for the parameters of the Correlated Chain Ladder (CCL). The higher posterior density regions (HPD) give an immediate overview of the learning update for each parameter. The thick line represents the 90% HPD while the thinner one extends it to a 95% HPD (Indexes start from 1).

With respect to the correlation parameter ρ , it is worth looking at its posterior distribution. In Figure 2, we see how the learning process affected this parameter. We see how the posterior is no longer uniform with a peak slightly higher than zero, thus implying that the posterior correlation is almost zero.




Figure 2. Posterior distribution of the correlation parameter. We see that the learning update implied a distribution no longer uniform and concentrated slightly above 0.

4.2. Main Results in a Run-Off View

After performing checks and studying the posterior distributions, we generate a predictive distribution for the claims reserve on a total run-off time frame and we analyze the results of the model. To this end, in Table 4, we compare for different accident years main characteristics of claims reserve computed by means of CCL and *Bootstrap-ODP*. Similarly, in Table 5, we provide the total claims reserve, compared with both the bootstrap ODP and the Mack formula. Finally, we plot the predictive distribution in Figure 3. With respect to the accident years, both methods show comparable results. CCL provides a higher mean and variability on the youngest accident years because of a broader prior. The total predictive distribution shows a mean of 205 million of euros, slightly lower than the results obtained with the Bootstrap-ODP and the Mack formula. The volatility returned by the CCL (for instance, CV is equal to 9.7%) is comparable to the one provided by the Bootstrap-ODP (CV equal to 9%). Both values are higher than the volatility provided by formula proposed by Mack. The most interesting comparison anyway is about the shape of the distribution where the CCL shows higher skewness, kurtosis (equal to 0.46 and 4.43 respectively), and upper quantiles than Bootstrap. In general, we observe that the CCL seems to better describe the usual heaviness of the right tail.

Table 4. We provide a comparison between CCL and Bootstrap-ODP method. For both methods, we report main characteristics (mean, coefficient of variation (CV), and 99% quantile). All the amounts are expressed in thousands of Euros and differentiated for accident years.

Years	Reserve		CV		99% Quantile	
	CCL	Bootstrap-ODP	CCL	Bootstrap-ODP	CCL	Bootstrap-ODP
1	144.21	502.86	4.27	1.61	1914.01	3148.01
2	910.77	1259.90	0.67	0.91	2681.07	4706.00
3	3138.16	3448.71	0.21	0.50	4901.30	8262.38
4	9010.50	8305.66	0.18	0.32	13,046.03	15,440.29
5	11,923.17	13,865.66	0.18	0.25	16,994.27	22,596.10
6	10,111.27	11,535.85	0.18	0.26	14,627.64	19,249.20
7	15,857.02	15,843.07	0.15	0.22	21,780.24	24,550.01
8	17,128.40	18,939.31	0.17	0.20	24,191.44	28,935.08
9	32,759.73	32,282.18	0.15	0.16	45,654.94	45,031.20
10	104,906.96	103,560.53	0.17	0.12	153,970.68	133,452.21

Table 5. Mean, standard deviation (SD), and CV of claims reserve obtained by CCL, Bootstrap-ODP, and Mack formula, respectively. For CCL and Bootstrap-ODP methods, skewness, kurtosis, and quantiles of the distribution are also reported. All the amounts are expressed in thousands of Euros.

	CCL	Bootstrap-ODP	Mack
Mean	205,890.19	209,543.74	209,255.94
SD	19,912.03	18,872.71	16,335.99
CV	0.097	0.090	0.078
Skewness	0.46	0.06	-
Kurtosis	4.43	3.05	-
1st quartile	192,706.64	196,582.00	-
Median	204,958.17	209,066.00	-
3rd quartile	217,790.79	222,285.75	-
99% quantile	257,935.78	254,001.40	-
99.5% quantile	268,426.73	259,138.41	-
Max	357,274.68	282,362.20	-




Figure 3. Predictive distribution (10,000 realizations) of claims reserve (CCL methodology).

4.3. One-Year View

We show the results of the model on the one-year time frame. The starting point is given by the parameters sets generated in the context of the total run-off analysis. We proceeded generating 10,000 samples of the one-year predictive distribution of claims reserve according to the described algorithm. In Table 6, with regard to the whole next-year obligations, we show the results of the one-year CCL alongside with the same results produced by the Re-Reserving method, the one-year frequentist simulative model for claims reserves. Similarly, we include the results produced by applying the Merz–Wüthrich formula to our triangle. The distribution produced by the CCL has a mean of 206 million euros and a coefficient of variation of 8.4%. It captures about 87% percent of the total run-off variability and still shows the classical features of claims reserves distributions with a not negligible skewness and some excess of kurtosis. We applied the CCL as an internal model by computing the solvency capital requirement (SCR) with a $Var_{99.5\%}$ risk measure obtained by subtracting the Best Estimate¹¹. The model turned out to provide an SCR equal to 52 million euros, equivalent to the 25.5% of the CCL Best Estimate. We then followed the same procedure for Re-Reserving. This model shows lower variability and skewness, thus implying a lower SCR (equal to 18.1% of the CL Best Estimate) mainly due to a weaker estimate of the right tail. This is also confirmed by a very different behavior in terms of Expected Shortfall. Finally, we compared the results of the simulative models with the Merz–Wüthrich closed formula fitted on a log-normal distribution. In this case, parameters of the log-normal distribution have been obtained using the CL Best Estimate and the σ_{usp}^{M-W} that the regulation allows to use as an *undertaking specific parameter*. As provided by the Solvency II regulation, the relative volatility has been derived weighting the CV (equal to 6.37%), obtained as a ratio between the MW standard error and the Best Estimate, with the *market wide* parameter. The weighting has been made using a *credibility factor* that took into consideration the time depth of the triangle. The SCR ratio resulted to be equal to 19.6%, a bit higher than the value obtained by the Re-Reserving methodology because of the higher skewness implied by the log-normal assumption.

Table 6. We provide all the main results obtained applying three alternative one-year models to our triangle. Figures in italics regarding the Merz–Wüthrich formula have been produced by assuming a log-normal distribution for the claims reserve fitted on the BE and the σ_{usp}^{M-W} that the regulations would have allowed to use as *undertaking specific parameter*. Amounts in thousands of euros.

	One-Year CCL	Re-Reserving	Merz–Wüthrich
Best Estimate	206,902.20	210,651.33	210,651.33
Mean of one-year obligations	205,756.09	209,184.98	209,255.94
SD	17,325.32	14,748.63	13,421.28
CV	0.084	0.071	0.071
One Year on Total	0.87	0.78	0.90
Skewness	0.36	0.18	0.21
Kurtosis	3.86	2.98	3.08
Quantile 75%	216,180.39	219,181.26	220,356.51
Quantile 95%	251,118.22	245,427.65	247,556.06
Quantile 99.5%	259,668.71	248,781.01	251,946.20
Max	313,101.79	267,166.99	-
Expected Shortfall 99%	262,529.18	250,566.40	253,601.99
Reserve Risk SCR	52,766.51	38,129.68	41,294.86
SCR over BE	25.50%	18.10%	19.60%

¹¹ For each parameters set we calculated the expected value of the lower triangle and discounted it according to the December 2018 Eiopa term structure. Subsequently, we obtained a CCL BE by calculating the mean of these discounted expected values.

To conclude, we notice that the one-year CCL is not only in line with the more traditional approaches, but can qualify as a model able to produce more robust estimates of the tail of our interest, and a third alternative approach when the other two do not provide a conclusive evidence. Finally, we provided in Figure 4 a plot of the one-year distribution generated by this adapted CCL.




Figure 4. Distribution of the claims reserve over a one-year time horizon, obtained from 10,000 simulations. The lines represent the 99.5% quantile and the Best Estimate, respectively. By subtracting the Best Estimate from the quantile, we obtain the Solvency II capital requirement that represents the hypothetical SCR for reserve risk returned by this internal model. As reported in Table 6, this value is equal to 52,766.51. Amounts are in thousands of euros.

4.4. Important Remarks

The one-year CCL proved to be computationally intensive. In fact, for all the 10,000 sets of parameters, we generated 10,000 batches of simulations of the first diagonal, each composed of other 10 simulations. The algorithm resulted in approximately 10^9 simulations to get 10,000 realizations relative to an 11×11 triangle. The method took a while to produce these results and it is advised both to implement it on a fast language as C and to exploit the possibilities of *parallel computing* to split the work on many different cores, or *distributed computing*, to split the work on different machines. In this study, we implemented the model in R, exploiting the possibilities offered by the package *parallel* that allowed us to split the calculations on three different cores of a laptop. The code to reproduce these results is available on the github profile 'GiulioErcole'.

5. Conclusions

At end of this study, the one-year CCL performed well and proved to be a legitimate alternative to the other well known models to quantify reserve risk. We sense that this model has something more to offer with respect to Re-Reserving, Merz–Wüthrich formula and the *market wide parameter* as it captures far more information about the variability of loss liabilities and provide a more truthful representation of the risks connected to the claims reserve. Despite the computational issues, the Bayesian nature of this model allows for embedding prior and external knowledge that could potentially lead to more

accurate and tailored representations of the risk profile of the insurer and that are even more in line with the principles behind the possibilities of structuring internal models. Thus we are convinced that more in general a Bayesian approach has still a lot of potential in light of the Solvency II framework. Anyway, it is to be kept in mind that this feature could be a double edged sword as this approach includes some sort of subjectivity that poses the risk that a company could distort purposefully their estimates in order to relieve its capital charges. As a result of this, we want to stress that it is extremely important that priors are elicited truthfully in order to obtain sound and plausible estimates.

Author Contributions: All authors contributed equally to this manuscript. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

Legal Disclaimer: The opinions expressed in the paper are solely those of the author. This paper is for general information, education and discussion purposes only. It does not constitute legal or professional advice and does not necessarily reflect, in whole or in part, any corporate position, opinion or view of PartnerRe or its affiliates.

Appendix A. An N-Years Generalization

The algorithm shown above could be easily generalized on an n -year time frame, allowing to obtain a predictive distribution for the *next-n-years obligations*:

$$\tilde{Y}^{n-yrs} = \sum_{h=1}^n \left(\sum_{i=h}^t P_{i,t-i+h} \right) + R^{D(t+n)} \tag{A1}$$

From this distribution it is possible to obtain a hypothetical capital charge on such a time frame by applying a risk measure. A longer time period can be useful for instance for own risk assessment purposes as provided by the second pillar of Solvency II. What follows is an extension of the Solvency II guidelines on an n -years time horizon:

1. Starting from the point 2 of the outline of the algorithm structured in the previous section, we simulate a $K \cdot S$ array in which every element is composed by a number n of diagonals, where again K is the number of parameter sets and S is the number of simulations for each parameter set. The s, k -th set of diagonals represent an hypothetical development of the triangle on which we will build a *recursive update*. In the one-year model, we had for the first diagonal:

$$P[\Theta^{(k)} | T_1^{(s,k)}] = \frac{L[T_1^{(s,k)} | \Theta^{(k)}] \cdot P(\Theta^{(k)})}{\sum_{h=1}^K L[T_1^{(s,h)} | \Theta^{(h)}] \cdot P(\Theta^{(h)})} = \frac{L[T_1^{(s,k)} | \Theta^{(k)}]}{\sum_{h=1}^K L[T_1^{(s,h)} | \Theta^{(h)}]}$$

where T_1 stands for the trapezoid obtained by adding 1 future diagonal to the triangle. We can then use the probabilities obtained in this way to update our knowledge after a second year of run-off:

$$P[\Theta^{(k)} | T_2^{(s,k)}] = \frac{L[T_2^{(s,k)} | \Theta^{(k)}] \cdot P[\Theta^{(k)} | T_1^{(s,k)}]}{\sum_{h=1}^K L[T_2^{(s,h)} | \Theta^{(h)}] \cdot P[\Theta^{(h)} | T_1^{(s,k)}]}$$

Then, recursively obtain the probabilities of each scenario after n years:

$$P[\Theta^{(k)} | T_n^{(s,k)}] = \frac{L[T_n^{(s,k)} | \Theta^{(k)}] \cdot P[\Theta^{(k)} | T_{n-1}^{(s,k)}]}{\sum_{h=1}^K L[T_n^{(s,h)} | \Theta^{(h)}] \cdot P[\Theta^{(h)} | T_{n-1}^{(s,k)}]} \tag{A2}$$

Iterating over S , we have S probability distributions that we will use as before to compute S determinations of the residual reserve after n years.

2. At this point, for each parameters' set we calculate the Best Estimate at time $t + n$ for claims incurred at time t . As for the one-year version, we compute the expected values of the lower triangle:

$$E[C_{ij}^{(k)}] = exp\left\{ \mu_{ij}^{(k)} + \frac{\sigma_j^{2(k)}}{2} \right\} \tag{A3}$$

for $i + j > t + n$. Again, from these values we can obtain the expected values of incremental amounts and the new Best Estimate for a given parameters' set $R_k^{D(t+n)}$.

- By iterating step 2 on all parameters sets we obtain K values of the *next-n-years* Best Estimate. For each batch of simulations s we use the posterior distribution of parameters' sets computed at step 1 in order to reweight the Best Estimates in light of the realized simulations. For every s we obtain one value:

$$\hat{R}_s^{D(t+n)} = \sum_{k=1}^K R_k^{D(t+n)} \cdot P[\Theta^{(k)} | T_n^{(s,k)}] \tag{A4}$$

- As before we use the simulated cumulative developments stored in the array to obtain a new array of simulated incremental developments:

$$\begin{aligned} P_{i,t-i+h}^{(s,k)} &= C_{i,t-i+h}^{(s,k)} - C_{i,t-i+h-1}^{(s,k)} \quad \text{for } h = 2, \dots, n \\ P_{i,t-i+1}^{(s,k)} &= C_{i,t-i+1}^{(s,k)} - C_{i,t-i} \end{aligned} \tag{A5}$$

for $i = 0, \dots, t$.

- Having at disposal simulated diagonals and the n years Best Estimates it is possible to obtain the predictive distribution of the *next-n-years* obligations. For every value of the Best Estimate we add a realized development of payments for the first n years. For the s -th value of the Best Estimate we have K simulated scenarios of developments. Hence, we sample a scenario and we add it to the s -th realization of the n years Best Estimate:

$$Y_s^{n-yrs} = \left(\sum_{h=1}^n \left(\sum_{i=h}^t P_{i,t-i+h}^{(s,w)} \right) \right) + \hat{R}_s^{D(t+n)} \tag{A6}$$

where w is a random number between 1 and K. Again, by iterating this process over all the S batches of simulations, we finally obtain a predictive distribution of the *next-n-years* obligations.

Clearly the methodology described in the n -years time frame is a generalization of the one-year approach reported in the previous section. This extension can be useful for risk margin purposes because it is possible to modify a bit Equation (A6) in order to catch the pattern of reserve risk capital requirement until total run-off. Additionally, analyses on a longer time frame are useful in the context of capital management and decision making. For instance, the main purpose of the forward looking assessment of the undertaking's own risks, is to ensure that the undertaking engages in the process of assessing all the risks inherent to its business and determines the corresponding capital needs. As specified in the Solvency II regulation, the results of the own risk solvency assessment process need to be developed including medium term capital management.

Appendix B. Data

Here, we provide the triangle and the premium that have been used as a starting point for the analyses contained in this paper. The triangle covers 11 years and is representative of a *motor third party liability* portfolio. To preserve confidentiality of the data, claims have been generated from a log-normal distribution.

Table A3. Expected value of the parameter μ of the a-posteriori distribution.

Accident Years	Development Years										
	0	1	2	3	4	5	6	7	8	9	10
0	10.95	11.60	11.70	11.74	11.77	11.78	11.81	11.83	11.85	11.85	11.85
1	11.04	11.70	11.80	11.84	11.87	11.88	11.90	11.93	11.94	11.94	
2	11.05	11.71	11.81	11.85	11.88	11.89	11.91	11.94	11.95		
3	11.07	11.73	11.83	11.87	11.90	11.91	11.93	11.96			
4	11.23	11.88	11.98	12.02	12.05	12.06	12.09				
5	11.32	11.98	12.08	12.12	12.15	12.16					
6	10.99	11.64	11.74	11.79	11.81						
7	11.03	11.69	11.79	11.83							
8	10.88	11.53	11.63								
9	10.98	11.64									
10	11.17										

Table A4. Comparison between loss ratio a priori and posterior value.

Years	Prior Loss Ratio	Posterior Loss Ratio
1	0.73756	0.7359
2	0.7178	0.7163
3	0.731	0.7282
4	0.806	0.8045
5	0.85	0.8487
6	0.685	0.6838
7	0.787	0.7863
8	0.705	0.7053
9	0.72	0.7209
10	0.74	0.7407

References

- Berger, James O., José M. Bernardo, and Dongchu Sun. 2009. The formal definition of reference priors. *The Annals of Statistics* 37: 905–38. [CrossRef]
- Bornhuetter, Ronald L., and Ronald E. Ferguson. 1972. The actuary and IBNR. In *Proceedings of the Casualty Actuarial Society*. Arlington: Casualty Actuarial Society, pp. 181–95.
- Dacorogna, Michel, Alessandro Ferriero, and David Krief. 2018. One-Year Change Methodologies for Fixed-Sum Insurance Contracts. *Risks* 6: 75. [CrossRef]
- Diers, Dorothea. 2009. Stochastic re-reserving in multi-year internal models—An approach based on simulations. Presented at Astin Colloquium, Helsinki, Finland, June 1–4.
- Efron, Bradley. 2011. The Bootstrap and Markov Chain Monte Carlo. *The Journal of Biopharmaceutical Statistics* 21: 1052–62. [CrossRef] [PubMed]
- England, Peter D., and Richard Verrall. 1999. Analytic and bootstrap estimates of prediction errors in claims reserving. *Insurance Mathematics and Economics* 25: 281–93. [CrossRef]
- European Commission. 2009. *Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the Taking-Up and Pursuit of the Business of Insurance and Reinsurance (Solvency II)*. Luxembourg: European Commission.
- Frees, Edward W., Richard A. Derrig, and Glenn Meyers, eds. 2016. *Predictive Analytics Applications in Actuarial Science*. Cambridge: Cambridge University Press.
- Friedland, Jacqueline. 2010. Estimating Unpaid Claims Using Basic Techniques. *Casualty Actuarial Society*. Available online: <https://www.casact.org/press/index.cfm?fa=viewArticle&articleID=816> (accessed on 19 November 2020).
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning, Data Mining Inference and Prediction*, 2nd ed. Stanford: Stanford University.
- Hindley, David. 2017. Deterministic Reserving Methods. In *Claims Reserving in General Insurance (International Series on Actuarial Science)*. Cambridge: Cambridge University Press, pp. 40–145.

- Hoffman, Matthew D., and Andrew Gelman. 2014. The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* 15: 1351–81.
- Klugman, Stuart A. 1992. *Bayesian Statistics in Actuarial Science*. Boston: Kluwer.
- Mack, Thomas. 1993. Distribution Free calculation of the standard error of Chain Ladder reserve estimates. *ASTIN Bulletin The Journal of the IAA* 23: 213–25. [[CrossRef](#)]
- Meyers, Glenn. 2017. A Cost of Capital Risk Margin Formula For Non-Life Insurance Liabilities. *Variance*.
- Meyers, Glenn. 2015. *Stochastic Loss Reserving Using Bayesian MCMC Models*. CAS monograph series. Arlington: Casualty Actuarial Society.
- Ntzoufras, Ioannis, and Petros Dellaportas. 2002. Bayesian Modelling of Outstanding Liabilities Incorporating Claim Count Uncertainty. *North American Actuarial Journal* 6: 113–28. [[CrossRef](#)]
- Peters, Gareth W., Rodrigo S. Targino, and Mario V. Wüthrich. 2017. Full bayesian analysis of claims reserving uncertainty. *Insurance Mathematics and Economics* 73: 41–53. [[CrossRef](#)]
- Wüthrich, Mario V., and Michael Merz. 2007. *Stochastic Claims Reserving Methods in Insurance*. Hoboken: Wiley.
- Wüthrich, Mario V., Michael Merz, Hans Bühlmann, Massimo De Felice, Alois Gisler, and Franco Moriconi. 2008. *Modelling the Claims Development Result for Solvency Purpose*. CAS E-Forum.
- Wüthrich, Mario V. 2007. Using a Bayesian Approach for Claims Reserving. *Variance* 1: 292–301.

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).